Submitted to the Annals of Statistics

arXiv:1608.05173v1 [stat.CO] 18 Aug 2016

APPROXIMATE BAYESIAN COMPUTATION VIA SUFFICIENT DIMENSION REDUCTION

BY XIAOLONG ZHONG
AND MALAY GHOSH *

University of Florida

Approximate Bayesian computation (ABC) has gained popularity in recent years owing to its easy implementation, nice interpretation and good performance. Its advantages are more visible when one encounters complex models where maximum likelihood estimation as well as Bayesian analysis via Markov chain Monte Carlo demand prohibitively large amount of time. This paper examines properties of ABC both from a theoretical as well as from a computational point of view. We consolidate the ABC theory by proving theorems related to its limiting behaviour. In particular, we consider partial posteriors, which serve as the first step towards approximating the full posteriors. Also, a new semi-automatic algorithm of ABC is proposed using sufficient dimension reduction (SDR) method. SDR has primarily surfaced in the frequentist literature. But we have demonstrated in this paper that it has connections with ABC as well.

1. Introduction. There are two main objectives of this article. First, we want to provide some theoretical results related to the currently emerging topic of approximate Bayesian computation (ABC). The second is to show some connectivity between ABC and another important emerging topic of research, namely, sufficient dimension reduction (SDR). While the latter has surfaced primarily in the frequentist's domain of research, it is possible to tie it with ABC as well. In particular, we want to show how ABC can be carried through nonlinear SDR.

Modern science invokes more and more Byzantine stochastic models, such as stochastic kinetic network (Wilkinson (2011)), differential equation system (Picchini (2014)) and multi-hierarchical model (Jasra et al. (2012)), whose computational complexity and intractability challenge the application of classical statistical inference. Traditional maximum likelihood methods will malfunction when the evaluation of likelihoods becomes slow and inaccurate. Lack of analytical form of the likelihood also undermines the usage of Bayesian inferential tools, such as Markov chain Monte Carlo (MCMC), Laplace approximation (Tierney and Kadane (1986)), variational Bayes (Jaakkola and Jordan (2000)) and posterior expansion

*Supported by an NSF grant
Primary 62F15; secondary 62E20
Keywords and phrases: Approximate Bayesian computation, Bernstein–von Mises theorem, Inferential model, Nonlinear sufficient dimension reduction

(Johnson (1970), Zhong and Ghosh).

The ABC methodology stems from the observation that the interpretability of the candidate model usually leads to an applicable sampler of data given parameters, and ingeniously circumvents the evaluation of likelihood functions. The idea behind ABC can be summarized as follows:

Algorithm 1 Idea of ABC

1 Sample parameters θ_i from the prior distribution $\pi(\theta)$;

2 Sample data Z_i based on the model $f(z \mid \theta_i)$;

3 Compare the simulated data Z_i and the observed data $X_{i,\mathrm{obs}}$, to accept or reject θ_i.

Rubin (1984) first mentioned this idea and Tavaré et al. (1997) proposed the first version of ABC, while studying population genetics. The prototype of ABC in recent research was given in Pritchard et al. (1999), where the comparison of two data sets was simplified to a comparison of summary statistics S and the accept-reject decision was made up to a certain error tolerance. We can view this algorithm

Algorithm 2 Prichard's Modified ABC

1 Sample parameters θ_i from the prior distribution $\pi(\theta)$;

2 Sample data Z_i based on the model $f(z \mid \theta_i)$;

3 Accept θ_i if $\rho(S(Z_i), S(X_{\mathrm{obs}})) \leq \varepsilon$, for some metric ρ.

as a modified version of accept-reject algorithm (Robert and Casella (2013)). The posterior is sampled by altering the frequency of the proposal distribution, that is, the prior. Now the full posterior distribution is approximated by the following two steps (Fearnhead and Prangle (2012)):

$$\pi(\theta \mid X_{\mathrm{obs}}) \approx \pi(\theta \mid S_{\mathrm{obs}}) \approx \pi(\theta \mid S_{\mathrm{sim}} \in O(S_{\mathrm{obs}}, \varepsilon)), \tag{1.1}$$

where $O(S_{\mathrm{obs}}, \varepsilon)$ means a neighborhood defined by the comparison measure ρ and tolerance level ε. We may note that the first approximation is exact when S is sufficient. Allowing the summary statistics to vary in an acceptable range sacrifices a little accuracy in exchange for a significant improvement in computational efficiency, which makes the algorithm more practical and user-friendly.

Pursuant to Algorithm 2, there are multiple generalizations in the statistical literature. Marjoram et al. (2003) introduced MCMC-ABC algorithm to concentrate the samples in high posterior probability region, thereby increasing the accept rate. Noisy ABC, proposed by Wilkinson (2013), makes use of all the prior samples by assigning kernel weights instead of hard-threshold accept-reject mechanism and hence reduces the computational burden. This perspective is corroborated in

Fearnhead and Prangle (2012) by convergence of Bayesian estimators. When the dependence structure between hierarchies is intractable, ABC filtering technique innovated by Jasra et al. (2012) comes to the rescue. Later in Dean et al. (2014), a consistency argument is established for the specific case of hidden Markov models. Moreover, many ABC algorithms above can be easily coded in a parallel way, and hence take advantages of modern computer structures. This feature makes ABC algorithms extremely time-saving in comparison with long-established, looping-based MCMC and MLE algorithms.

Despite the fruitful results on ABC both from applied and theoretical points of view, there exist only a handful of papers which focus on the effect of the choice of summary statistics on the approximation quality. The quintessential case is when the summary statistics are sufficient, and the resultant ABC sampler produces exact samples from the true posterior distribution when ε goes to zero. Nevertheless, in a labyrinthine model, it is difficult to extract sufficient statistics, except for some very special cases, such as exponential random graph models (e.g. Grelaud et al. (2009)). Joyce and Marjoram (2008) proposed a concept called $\varepsilon-$sufficient to quantify the effect of statistics. Nonetheless, this property is also difficult to verify in complicated models. If we are interested only in model selection, Prangle et al. (2014) designed a semi-automatic algorithm to construct summary statistics via logistic regression. And laterly, Marin et al. (2014) gave sufficient conditions on summary statistics in order to choose the right model based on the Bayes factors. They advocated that the ideal summary statistics are ancillary in both model candidates. One of our contribution comes from the mathematical analysis of the consequence of conditioning the parameters of interest on consistent statistics and intrinsically inconsistent statistics, and appraises the efficiency of the posterior approximation based on the former. Generally speaking, using consistent statistics results in right concentration of the approximate posterior, while less efficient statistics lead to less efficiency of approximation. One byproduct is our theorem vindicates the usage of the posterior mean as summary statistics as in Fearnhead and Prangle (2012).

In addition to the pure theoretical contribution, we also extend the two-step algorithm in Fearnhead and Prangle (2012) in a more flexible and nonparametric way for automatic construction of summary statistics. We borrow the idea from another thriving topic, namely sufficient dimension reduction (SDR). The motivation of SDR which generalizes the concept of sufficient statistics is to estimate a transformation φ, either linear or nonlinear, such that

$$Y \perp\!\!\!\perp X \mid \varphi(X)\,. \tag{1.2}$$

The first SDR method titled sliced inverse regression dates back to Li (1991), followed by principle Hessian direction in Li (1992) and also by Cook and Weisberg (1991) and Cook (1998). As we step in the era of big data, this idea leads to a

sea of papers on both linear and nonlinear predictors and response. Among the more recent work, we refer to Cook and Li (2002), Xia et al. (2002), Li, Zha and Chiaromonte (2005), Li and Dong (2009), Wu (2008), Yeh, Huang and Lee (2009), Su and Cook (2011) and Su and Cook (2012). The association between SDR and ABC relies on the shared mathematical formulation. If we think θ as the response and X as the predictor, then an ideal summary statistics $S(X)$ will give

$$\theta \perp\!\!\!\perp X \mid S(X).$$

This simple observation offers raison d'etre to use existing SDR methods in constructing summary statistics. The employment of dimension reduction methods in our algorithm is different from that in Blum et al. (2013). In Blum et al. (2013), dimension reduction methods, such as best subset selection, projection techniques and regularization approaches, are applied to reduce the dimension of existing summary statistics, but here, we try to reduce the size of the original data. Particularly in our paper, we incorporate the principal support vector machine for nonlinear dimension reduction given in Li, Artemiou and Li (2011) into ABC, which uses the principal component of support vectors in reproducing kernel Hilbert space (RKHS) as a nonparametric estimator of φ.

The outline of remaining sections is as follows. Section 2 contains asymptotic results on the partial posterior. We gradually relax the restriction on summary statistics and investigate the relationship between the partial posterior and the full posterior. As a side result, we give a lemma building a bridge between the recent prior free inferential model (Martin and Liu (2013), Martin and Liu (2015)) and traditional Bayesian inference. Section 3 elicits a new ABC algorithm which automatically produces summary statistics through nonlinear SDR. A simulation result is provided in this section as well. Section 4 briefly discusses the results and points out some possible future generalizations.

2. Asymptotic Properties of Partial Posterior. Suppose $X_1, \ldots, X_n \mid \theta$ are i.i.d. with common PDF $f(x \mid \theta)$, and there exists a true but unknown value θ_0. Without loss of generality, we assume $\theta \in \mathbb{R}$, and all probability density functions are with respect to the Lebesgue measure. For illustration purpose, we define the following terminology.

DEFINTITION 1 (Partial Posterior). Let $S = S(X_1, \ldots, X_n)$ be statistics of the data. Given a prior $\pi(\theta)$, we call the distribution

$$\pi(\theta \mid S) \propto \pi(\theta)\, g(S \mid \theta)$$

the partial posterior, where $g(S \mid \theta)$ is the probability density function of statistic $S(X_1, \ldots, X_n)$ derived from the data density, and correspondingly,

$$\pi(\theta \mid X_1, \ldots, X_n) \propto \pi(\theta)\, f(X_1, \ldots, X_n \mid \theta)$$

is called the full posterior.

From equation (1.1), the partial posterior significantly reduces the complexity of the full posterior by replacing the dependence on full data by lower dimensional statistics S. If the partial posterior deviates from the full posterior too much, then no matter how delicately we sample from $\pi\left(\theta \mid S_{\text{sim}} \in O\left(S_{\text{obs}}, \varepsilon\right)\right)$, and how small ε we choose, the resultant samples would not behave like ones drawn from the original full posterior. This makes the subsequent Bayesian analysis fragile and unreliable. Therefore, theoretical connection between some easily verifiable properties and asymptotic behaviour of the partial posterior is of relevance. In particular, we want to study consistency and asymptotic normality of our Bayesian procedures. The following theorems try to demonstrate the connection between the asymptotic behaviour of summary statistics and that of partial posterior. We start from the most popular statistics, the maximum likelihood estimators (MLE) of θ.

THEOREM 1. *Let $\hat{\theta}$, the MLE of θ, be a strongly consistent estimator, and let $\hat{I}$ be the observed Fisher information evaluated at $\hat{\theta}$, and the full posterior satisfies the Bernstein–von Mises theorem. Then for any $\varepsilon > 0$, and any t, the partial posterior after conditioned on $\hat{\theta}$ satisfies*

$$\lim_{n\to\infty} \operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid \hat{\theta} \in O\left(\theta_0, \varepsilon\right)\right\} = \Phi\left(t\right), \text{ a.s.}.$$

PROOF. See Appendix A. □

REMARK 1. There is a slight difference between

$$\lim_{n\to\infty} \operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid \hat{\theta} \in O\left(\theta_0, \varepsilon\right)\right\} = \Phi\left(t\right), \text{ a.s.}.\left(P_{\theta_0}\right)$$

and

$$\lim_{n\to\infty} \operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid \hat{\theta}\right\} = \Phi\left(t\right), \text{ a.s.}.$$

By definition,

(2.1)

$$\operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid \hat{\theta}\right\} = \lim_{\varepsilon\to 0} \frac{\operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t, \hat{\theta} \in O\left(s, \varepsilon\right)\right\}}{\operatorname{pr}\left\{\hat{\theta} \in O\left(s, \varepsilon\right)\right\}}.$$

The result of Theorem 1 can only be used to prove

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty} \operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid \hat{\theta} \in O\left(\theta_0, \varepsilon\right)\right\} = \Phi\left(t\right), \text{ a.s.},$$

switching order of limits in equation (2.1).

REMARK 2. The definition of $\mathrm{pr}\left\{\theta \mid \hat{\theta} \in O\left(\theta_0, \varepsilon\right)\right\}$ is different from the approximation $\mathrm{pr}\left\{\theta \mid \hat{\theta} \in O\left(\hat{\theta}_{\mathrm{obs}}, \varepsilon\right)\right\}$. In first case, $\hat{\theta}$ is evaluated at $X_1, \ldots, X_n \sim f\left(x \mid \theta_0\right)$, the observed data, while the latter evaluates $\hat{\theta}$ at $Z_1, \ldots, Z_m \sim f\left(z \mid \theta\right)$, the simulated data.

By assumptions , the asymptotic distribution of the full posterior is still normal, and we have

$$\begin{aligned} & \sup_{t \in \mathbb{R}} \left| \mathrm{pr}\left\{ \left(n\hat{I}\right)^{1/2} \left(\theta - \hat{\theta}\right) \leq t \mid \hat{\theta} \in O\left(\theta_0, \varepsilon\right) \right\} \right. \\ & \left. -\mathrm{pr}\left\{ \left(n\hat{I}\right)^{1/2} \left(\theta - \hat{\theta}\right) \leq t \mid X_1, \ldots, X_n \right\} \right| \\ \leq & \sup_{t \in \mathbb{R}} \left| \mathrm{pr}\left\{ \left(n\hat{I}\right)^{1/2} \left(\theta - \hat{\theta}\right) \leq t \mid \hat{\theta} \in O\left(\theta_0, \varepsilon\right) \right\} - \Phi\left(t\right) \right| \\ & + \sup_{s \in \mathbb{R}} \left| \mathrm{pr}\left\{ \left(n\hat{I}\right)^{1/2} \left(\theta - \hat{\theta}\right) \leq s \mid X_1, \ldots, X_n \right\} - \Phi\left(s\right) \right| \to 0, \ \left(\text{as } n \to \infty\right). \end{aligned}$$

Hence, we can informally say that two random variables $\left(n\hat{I}\right)^{1/2} \left(\theta - \hat{\theta}\right) \mid \hat{\theta}$ and $\left(n\hat{I}\right)^{1/2} \left(\theta - \hat{\theta}\right) \mid X_1, \ldots, X_n$ are close in distribution. Note that both random variables asymptotically center at consistent MLE, and hence will eventually concentrate at θ_0. Meanwhile, the scale factors in both random variables are $\left(n\hat{I}\right)^{1/2}$, which ensures the same square root credible intervals. In this sense, we feel that the partial posterior conditioned on the MLE has the same efficiency as the full posterior. Later theorems will tell us that if the summary statistics are not efficient, the corresponding partial likelihood will have a different scale factor, and thus will lose efficiency and result in a larger credible interval.

A slightly modified proof of Theorem 1 can be used to support the posterior mean as a summary statistic in Fearnhead and Prangle (2012) and we still have a similar result, namely

$$\begin{aligned} & \lim_{n \to \infty} \mathrm{pr}\left[\left(n\hat{I}\right)^{1/2} \left\{\theta - E\left(\theta \mid X_1, \ldots, X_n\right)\right\} \leq t \right. \\ & \left. \mid E\left(\theta \mid X_1, \ldots, X_n\right) \in O\left(\theta_0, \varepsilon\right) \right] = \Phi\left(t\right), \text{ a.s.}. \end{aligned}$$

The key fact to support the assertion above comes from Ghosh and Liu (2011), that

is, the higher order closeness of the posterior mean and the MLE, namely

$$\lim_{n\to\infty} n^{1/2}\left\{E\left(\theta \mid X_1,\ldots,X_n\right)-\hat{\theta}\right\}=0,\ \text{a.s.}. \tag{2.2}$$

Indeed, any estimator who has the same or higher order of closeness to MLE will work as an efficient summary statistic.

Theorem 1 can be generalized to more intricate models. The following example shows the same phenomenon in data generated from a Markov process.

EXAMPLE 1. Immigration-emigration process is a crucial model in survival analysis and can be viewed as a special case of mass-action stochastic kinetic network (Wilkinson (2011)). The model is defined by a birth procedure and a death procedure during an infinitesimal time interval, namely,

$$\text{pr}\left\{X\left(t+\ \mathrm{d}t\right)=x_1 \mid X\left(t\right)=x_0\right\}=\begin{cases}\lambda\,\mathrm{d}t+o\left(\mathrm{d}t\right), & x_1=x_0+1,\\ \mu x_0\,\mathrm{d}t+o\left(\mathrm{d}t\right), & x_1=x_0-1,\\ 1-\lambda\,\mathrm{d}t-\mu x_0\,\mathrm{d}t+o\left(\mathrm{d}t\right), & x_1=x_0.\end{cases}$$

Assume that we observe full data in the time interval $[0,T]$. Let $T_i, i=1,\ldots,n$ be the event times and let $X_i=X\left(T_i\right), i=1,\ldots,n$. Let X_0 be initial population, $T_0=0$, $T_{n+1}=T$. Then by Gillespie's algorithm, the likelihood is proportional to

$$\lambda^{r_1}\exp\left(-\lambda T\right)\mu^{r_2}\exp\left(-\mu A_T\right),$$

where r_1 and r_2 are number of events corresponding to immigration and emigration, and

$$A_T=\int_0^T X\left(t\right)\,\mathrm{d}t.$$

The MLEs are

$$\hat{\lambda}=\frac{r_1}{T},\hat{\mu}=\frac{r_2}{A_T},$$

and they are strongly consistent estimators of λ and μ when T goes to infinity. By the computation in Appendix B.1, we have the partial posterior density function of $T^{1/2}\left(\mu-\hat{\mu}\right)$ conditioned on $\hat{\mu}$, r_1 and T given by

$$\lim_{T\to\infty}\pi\left\{T^{1/2}\left(\mu-\hat{\mu}\right)=t \mid \hat{\mu},r_1,T\right\}=\frac{\hat{\mu}}{\left(2\pi\hat{\lambda}\right)^{1/2}}\exp\left(-\frac{\hat{\lambda}}{\hat{\mu}^2}t^2\right),\ \text{a.s.}.$$

The MLE seems to be a perfect surrogate for the full data. However, in many cases, use of MLE is prohibitive due to heavy computational burden, particularly when the likelihood function is intractable. This is when the ABC comes on stage.

M-estimator is a generalization of the MLE, which is also consistent and asymptotically normal under mild conditions. Many M-estimators can be easily calculated, especially some moment estimators. To give an idea of the nature of approximation, we consider the following examples.

EXAMPLE 2. Gamma distribution can be used to model hazard functions in survival analysis. The shape parameter of gamma distribution determines the trend of hazard and hence is a vital parameter to estimate. Assume $X_1, \ldots, X_n \sim \text{Gamma}(\alpha, \beta)$, where we know the scale parameter β, but not the shape parameter α. The MLE of α is the solution of

$$-\log\Gamma(\alpha) - \alpha\log\beta + (\alpha-1)\sum_{i=1}^{n}\log X_i - \frac{\sum_{i=1}^{n} X_i}{\beta} = 0,$$

which involves repeated evaluation of the gamma function in search of the root. A simple M-estimator $\tilde{\alpha} = \overline{X}/\beta$ is derived from its mean equation,

$$\sum_{i=1}^{n}(X_i - \alpha\beta) = 0.$$

Now we consider the partial posterior $\pi(\alpha \mid \tilde{\alpha})$, when the prior is $\pi(\alpha) \propto \exp(-\lambda\alpha)$. By the calculation in Appendix B.2, we show that the limit of cumulative probability function of $n^{1/2}\tilde{\alpha}^{-1}(\alpha - \tilde{\alpha})$ given $\tilde{\alpha}$ is

$$\lim_{n\to\infty} \text{pr}\left\{n^{1/2}\tilde{\alpha}^{-1}(\alpha-\tilde{\alpha}) \le t \mid \tilde{\alpha}\right\} = \Phi(t), \text{ a.s.},$$

which means that the Bernstein-von Mises theorem holds for the partial posterior conditioned on the M-estimator $\tilde{\alpha}$. The scale factor of the partial posterior is $n^{1/2}\tilde{\alpha}^{-1}$, which is smaller than that of the full posterior, $\{n\psi'(\alpha)\}^{1/2}$, where $\psi(\alpha)$ is digamma function. That results in a larger credible interval based on the partial posterior.

EXAMPLE 3. Another example is the Laplace distribution with pdf

$$f_{\mu,\lambda}(t) = \frac{1}{2\lambda}\exp\left(-\frac{|t-\mu|}{\lambda}\right).$$

Here we want inference for the location parameter μ holding λ fixed. The MLE is the the sample median and the moment estimator is the sample mean. Here we calculate the partial posterior based on the sample mean. By the calculation in Appendix B.3, we find that the characteristic function of $n^{1/2}\left(\mu - \overline{X}\right)$ converges to $\exp\left(-\lambda^2 t^2\right)$, which is the characteristic function of normal distribution.

Example 3 uses the following lemma which is of independent interest.

LEMMA 1. *Assume X has the same distribution as $h(Y,\theta)$, where $h(y,\theta)$ for a fixed y is a one-to-one function of θ and Y is a random variable independent of θ. Let $\theta = g(y,x)$ and $y = u(x,\theta)$ be the solutions of the equation $x = h(y,\theta)$. Further assume $\partial u(x,\theta)/\partial x$ exists and is not equal to zero. Then the posterior distribution of θ conditioned on X under the uniform prior has the same distribution as $g(Y,x)$, where x is fixed.*

REMARK 3. Although not quite related to ABC, this lemma gives another interpretation of inferential model of Martin and Liu (2013) and Martin and Liu (2015). In their settings, Y is called unobserved ancillary variable, and $g(y,x)$ is $\Theta_x(u)$ in their notation. They claim that their procedure results in a distribution of θ without referring to a prior. However, by our lemma, this model is mathematically the same as a posterior given a uniform prior.

The following theorems are built upon the Theorem 2.1 in Rivoirard et al. (2012), which guarantees the asymptotic normality of linear functionals of nonparametric posterior. So we need all the assumptions in that theorem. Additionally, we need the following assumptions.

ASSUMPTION 1. There is a neighbourhood $\theta \in O(\theta_0,\varepsilon)$ such that $\int_{\mathbb{R}} g(x,\theta_0)\,\pi(x \mid \theta)\,\mathrm{d}x$ is a continuous twice differentiable in θ and the second order derivative is bounded by some constant L.

ASSUMPTION 2. M-estimator $\tilde{\theta}$ and MLE $\hat{\theta}$ are both strongly consistent and asymptotically normal.

ASSUMPTION 3. Bernstein–von Mises theorem and posterior consistency hold for the full posterior of θ.

ASSUMPTION 4. For any $\theta \in \Theta$, $E_{\theta_0} \log f(X \mid \theta) \le E_{\theta_0} \log f(X \mid \theta_0)$.

Now we can articulate the theorem.

THEOREM 2. *Under the Assumptions 1, 2, 3, 4, and conditions of Theorem 2.1 in Rivoirard et al. (2012), for any ε and t,*

$$\lim_{n\to\infty} \mathrm{pr}\left\{ \left(n/\tilde{V}\right)^{1/2} \left(\theta - \tilde{\theta}\right) \le t \,\middle|\, \tilde{\theta} \in O(\theta_0,\varepsilon) \right\} = \Phi(t), \text{ a.s.},$$

where $\tilde{V} = V_0/G_1\left(\tilde{\theta},\tilde{\theta}\right)^2$ is the Godambe information.

PROOF. See Appendix C. □

Using similar arguments as Theorem 1, the partial posterior $\left(n/\tilde{V}\right)^{1/2}\left(\theta-\tilde{\theta}\right) \mid \tilde{\theta}$ is asymptotically close in distribution to the full posterior $\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \mid X_1,\ldots,X_n$. Since both the M-estimator and the MLE are strongly consistent, the partial posterior still concentrates around the right θ_0, but now the asymptotic $\alpha-$level credible interval based on the partial posterior, namely

$$\left(\tilde{\theta}-Z_{\alpha/2}\left(\frac{\tilde{V}}{n}\right)^{1/2},\tilde{\theta}+Z_{\alpha/2}\left(\frac{\tilde{V}}{n}\right)^{1/2}\right),$$

will be larger than that based on the full posterior,

$$\left(\hat{\theta}-Z_{\alpha/2}\left(\frac{\hat{I}^{-1}}{n}\right)^{1/2},\hat{\theta}+Z_{\alpha/2}\left(\frac{\hat{I}^{-1}}{n}\right)^{1/2}\right),$$

where $Z_{\alpha/2}$ is the $(1-\alpha/2)$ quantile of standard normal distribution. This is because the Godambe information $\tilde{V}^{-1}$ is typically no larger than Fisher information $\hat{I}$. Hence, we lose efficiency if we condition the posterior on an inefficient estimator, which coincides with our intuition.

For extreme tortuous models, even finding a consistent estimator can be quite hard. There are still some simple statistics which may be consistent to some functions of θ. Unless they are ancillary statistics, they always contain some information about the parameters of interest. Moreover, in the real case, we use several statistics, each of which gives independent information of the full posterior. In the remainder of this section, we will mathematically quantify what independent information means and show that using more than one statistic will improve the efficiency.

Let S_i, $i=1,\ldots,q$ be statistics of the sample. We make the following trivial assumptions.

ASSUMPTION 5. The joint distribution of $S_1,\ldots,S_q$ converges in distribution to a multivariate normal distribution $N\left(h\left(\theta_0\right),n^{-1/2}\Sigma\left(\theta_0\right)\right)$, and each S_i converges to $h_i\left(\theta_0\right)$ almost surely. Further, assume $\Sigma\left(\theta_0\right)$ is positive definite, and $h\left(\theta_0\right)$ is a linear functional of the distribution function, that is

$$h\left(\theta_0\right)=\int_{\mathbb{R}} g\left(x\right) f\left(x \mid \theta_0\right)\,\mathrm{d}x,$$

where $g\left(x\right)\in\mathbb{R}^q$.

Assumption 5 characterizes the independent information statement. Because if $\Sigma(\theta_0)$ has a lower rank, then some of S_i can be expressed as linear combinations of other S_j asymptotically. Then the partial posterior can be reduced to a partial posterior based solely on the S_j. The functional form of h is a natural consequence when we apply some version of strong law of large numbers to prove convergence of statistics.

In order to prove the theorem, we need some more technical assumptions.

ASSUMPTION 6. Let $S = (S_1, \ldots, S_q)$, assume

$$\lim_{n\to\infty} n^{1/2}\left\{\frac{1}{n}\sum_{i=1}^{n} g(X_i) - S\right\} = 0, \text{ a.s.}.$$

and there exists a strongly consistent estimator $\tilde{\Sigma}$ of $\Sigma(\theta_0)$

Only Assumption 6 seems quite restrictive. Based on all these assumptions, the theorem describing the partial posterior conditioned on less informative statistics can be found as follows:

THEOREM 3. *Under Assumptions 5, 6 and conditions of Theorem 2.1 in Rivoirard et al. (2012), for any vector* $a \in \mathbb{R}^q$,

$$\lim_{n\to\infty}\sup_{t\in\mathbb{R}}\left|\text{pr}\left[\frac{n^{1/2}a^T\{h(\theta) - S\}}{\left(a^T\tilde{\Sigma}a\right)^{1/2}} \le t \middle| S\right] - \Phi(t)\right| = 0, \text{ a.s.}.$$

PROOF. See Appendix D. □

Theorem 3 extends the asymptotic results about M-estimators to more general statistics, particularly the intrinsically inconsistent statistics defined as follows.

DEFINTITION 2 (Intrinsic Consistency). Let S be an non-ancillary statistic and converges to $h(\theta_0)$ almost surely. If $h(\cdot)$ is an one-to-one function and has an inverse function, then we say S is intrinsically consistent. Otherwise, we say S is intrinsically inconsistent.

If S is a one dimensional intrinsically inconsistent statistic, then Theorem 3 asserts the $(1-\alpha)$ asymptotic credible set based on the partial posterior is

$$\left\{\theta : S - Z_{\alpha/2}\left(\frac{\tilde{\Sigma}}{n}\right)^{1/2} \le h(\theta) \le S + Z_{\alpha/2}\left(\frac{\tilde{\Sigma}}{n}\right)^{1/2}\right\}.$$

In an extreme case, when sample size n is large enough, such that $Z_{\alpha/2}/\sqrt{n} \approx 0$, the asymptotic credible interval based on the the full posterior would be close to the singleton $\left\{\hat{\theta}\right\}$. However, the credible set based on the partial posterior would be $\{\theta : h(\theta) = S\}$. By the definition of intrinsic inconsistency, h is not a one-to-one function. Then the set $\{\theta : h(\theta) = S\}$ would possibly hold multiple elements, hence larger than that from the full posterior. Again, in this sense, we perceive loss of efficiency due to conditioning the posterior on arbitrary statistics.

Another interesting use of Theorem 3 is a more pragmatic asymptotic assessment of effectiveness of including many statistics than that in Joyce and Marjoram (2008). In their settings, the effectiveness of summary statistics is measured by the difference between log-likelihoods, thus not operable when likelihood functions are intractable. On the other hand, our approach only requires the asymptotic behaviour of statistics, and the corresponding credible set with q statistics can be developed by the Cramer device as

$$\left\{\theta : n\left(h(\theta) - S\right)^T \tilde{\Sigma}\left(h(\theta) - S\right) \leq \chi^2_{1-\alpha,q}\right\},$$

where $\chi^2_{1-\alpha,q}$ is $(1-\alpha)$ quantile of chi-square distribution with degree of freedom q. To select summary statistics, we can compare the asymptotic credible sets with and without the current statistic. If the difference is small, then we can safely throw the current statistic away.

The following is a simple example to illustrate this phenomenon.

EXAMPLE 4. Let X_i are i.i.d. sample from $N\left(\mu, \mu^2\right)$ and we calculate the partial posterior $\pi\left(\mu \mid s^2\right)$, where s^2 is the sample standard deviation. The prior of μ^2 is assumed to be inverse gamma distribution with shape parameter α and scale parameter β. We know the $s^2 \sim \mu^2\chi^2_{n-1}/(n-1)$, hence the partial posterior of μ^2 conditioning on s^2 is an inverse gamma distribution with shape parameter $\{\alpha - 1 + (n-1)/2\}$ and scale parameter $\left\{\beta + (n-1)s^2/2\right\}$. This leads to a bimodal partial posterior for μ. Hence, we can only get the absolute value of μ without sign information from this partial posterior.

3. Approximate Bayesian Computation via Nonlinear Sufficient Dimension Reduction. In principle, almost all the existing dimension reduction methods are valid in estimating the summary statistics. However, there is a slight difference between the setting of SDR and ABC. In the theory of SDR, the independent assumption 1.2 must hold rigorously, which implies $Y \mid X$ has exactly the same distribution as $Y \mid S(X)$. However, by our Theorem 1, 2 and 3, the two distributions are only close in large but finite samples.

Algorithm 3 Principal Support Vector Machine

1 (Optional) Marginally standardize data $X_1, \ldots, X_n$. The purpose of this step is so that the kernel κ treats different components of X_i more or less equally.

2 Choose kernel κ and the number of basis functions k (usually around $n/3 \sim 2n/3$). Compute $K = \{\kappa(X_i, X_j)\}_{n\times n}$. Let $Q = I_n - J_n/n$, where J_n is the $n \times n$ matrix whose entries are 1. Compute largest k eigenvalues $\lambda_1, \ldots, \lambda_k$ and corresponding eigenvectors $w_1, \ldots, w_k$ of matrix QKQ. Let $\Psi = (w_1, \ldots, w_k)$ and $P_\Psi = \Psi\left(\Psi^T\Psi\right)^{-1}\Psi^T$ be the corresponding projection matrix.

3 Partition the response variable space Y into h slices defined by $y_1, \ldots, y_{h-1}$. For each y_s, $s = 1, \ldots, h-1$, define a new response variable $\tilde{Y}_{si} = I_{[Y_i \le y_s]} - I_{[Y_i > y_s]}$. Then solve the modified support vector machine problem as a standard quadratic programming

$$\min_\alpha -1^T\alpha + \frac{1}{4}\alpha^T \mathrm{diag}\left(\tilde{Y}_s\right) P_\Psi \mathrm{diag}\left(\tilde{Y}_s\right)\alpha,$$

subject to constraints

$$\begin{cases} 0 \le \alpha \le \lambda, \\ \tilde{Y}_s^T\alpha = 0, \end{cases}$$

where $\mathrm{diag}\left(\tilde{Y}_s\right)$ is a diagonal matrix using $\tilde{Y}_s$ as diagonal, λ is a hyper-parameter in ordinary support vector machine. The coefficients of support vectors in RKHS are

$$c_s^* = \frac{1}{2}\left(\Psi^T\Psi\right)^{-1}\Psi^T \mathrm{diag}\left(\tilde{Y}_s\right)\alpha_s.$$

4 Let d be the target dimension. Compute the eigenvectors $v_1, \ldots, v_d$ of first largest d eigenvalues of the matrix $\sum_{s=1}^{h-1} c_s^* c_s^{*T}$. Let $V = (v_1, \ldots, v_d)$.

5 Let $K(x, X) = \left\{\kappa(x, X_i) - n^{-1}\sum_{j=1}^n \kappa(x, X_j)\right\}$ be a n dimensional vector. Then the estimated transformation $\hat{\varphi}(x) = V^T\left\{\mathrm{diag}(\lambda_1, \ldots \lambda_k)\right\}^{-1}\Psi^T K(x, X)$.

3.1. *Algorithm: ABC via PSVM.* In our paper, we choose principal support vector machine in Li, Artemiou and Li (2011). Suppose we have a regression problem (Y_i, X_i), and search a nonlinear transformation $\varphi : \mathbb{R}^p \to \mathbb{R}^d$, such that $Y \perp\!\!\!\perp X \mid \varphi(X)$. Then the main steps in principal support vector machine are given in Algorithm 3 .

By slicing the response variable space, we discretize variation of Y. The support vector machine in Step 3 recognizes the robust separate hyperplanes. We will expect the variation of Y along the directions within hyperplanes to be negligible and that along the directions perpendicular to the hyperplanes to explain the most part of covariation between Y and X. The principal component analysis on the support vectors in Step 4 estimates the principal perpendicular directions and hence creates the sufficient directions in RKHS.

Based on Algorithm 3, we formulate our two-step approximate Bayesian computation algorithm in Algorithm 4.

Algorithm 4 ABC via PSVM

1 Sample θ_i from the prior $\pi(\theta)$ and sample $X_{i1}, \ldots, X_{in}$ from the model $f(x \mid \theta_i)$.

2 View $(\theta_i, X_{i1}, \ldots X_{in})$ as a multivariate regression problem and reduce the dimension from n to d via principal support vector machine. Denote the estimated transformation as $\hat{S}(X_1, \ldots, X_n)$.

3 Either use existent samples in Step 1 or repeat it and get new sample. Calculate the estimated summary statistics $\hat{S}_i = \hat{S}(X_{i1}, \ldots, X_{in})$ on each set $X_{i1}, \ldots, X_{in}$ corresponding to prior samples θ_i. Also calculate $\hat{S}_{\text{obs}} = \hat{S}(X_1, \ldots, X_n)$ on the observed data set.

4 Based on the metric $\rho\left(\hat{S}_i, \hat{S}_{\text{obs}}\right)$, make the decision of accept or reject of θ_i.

Algorithm 4 directly generalizes the semi-automatic ABC in Fearnhead and Prangle (2012). In their algorithm, the summary statistics are fixed as posterior means and the recommended estimation method is polynomial regression. Our algorithm relaxes the restriction on summary statistics and lets the data and nonparametric algorithm together find them adaptively. One significant difference between our algorithm and the conventional ABC is in Step 1, where each prior sample θ_i produces exact n simulated data, because the nonparametric estimator of statistics should take n arguments so that it can be evaluated at both observed data and simulated data.

3.2. *Robustness of PSVM.* The discrepancy between the asymptotic behavior of ABC and settings in SDR requires new properties of PSVM, namely robustness, which suggests that $\hat{\Gamma}$ from PSVM would be in the vicinity of the true summary statistics function Γ in some sense even if the partial posterior is only close to the full posterior. Here is the theorem which validates this property.

THEOREM 4 (Robustness of PSVM). *Assume*

$$\lim_{n\to\infty} \sup_s \left|\text{pr}\left\{\theta \le s \mid \Gamma_n^T(X_1, \ldots, X_n)\right\} - \text{pr}(\theta \le s \mid X_1, \ldots, X_n)\right| = 0,$$

where $\Gamma \in \mathbb{R}^{n\times d}$, with d fixed. Further assume all the conditions in Theorem 4 and 5 in Li, Artemiou and Li (2011). Let $\hat{\Gamma}_n$ be the result from PSVM. Then

$$\hat{\Gamma}_n - \Gamma_n \xrightarrow{p} 0.$$

PROOF. See Appendix E. □

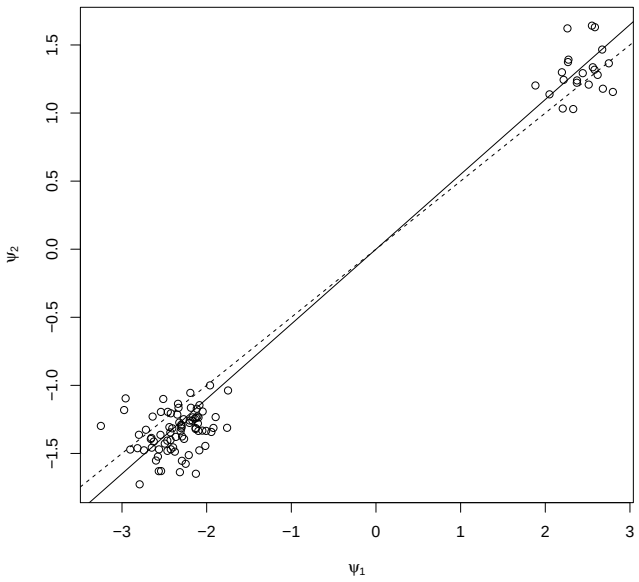

FIG 1. *Normal Vectors of Separate Hyperplanes in Sliced SVM*

3.3. *Simulation Example* . First, we will show a simple simulation example to illustrate the robustness of SVM, which serves as the foundation of the robustness of PSVM.

EXAMPLE 5. Let $\theta \mid X_1, X_2 \sim N\left(2X_1 + X_2 + 0.001\left(X_1^2 + X_2^2\right), 1\right)$. Hence the conditional distribution of θ on X_1, X_2 is close to $N\left(2X_1 + X_2, 1\right)$. Then the normal vectors of sliced SVM should be in vicinity of $(2, 1)$. We choose the slicing point $s = 1.5$. The results are summarized in Fig 1. The dots are normal vectors, the solid line is the principal component direction of normal vectors, and the reference dashed line is $\psi_1 = 2\psi_2$. We can see the principal component direction is quite close to the reference line.

Next, we will show a simple simulation example to illustrate our algorithm.

EXAMPLE 6. Autoregressive model with lag one, AR(1).

$$Y_i = \beta Y_{i-1} + \varepsilon.$$

Set $Y_1 = 1$ and number of observation is 100. Assume $\varepsilon \sim N\left(0, 0.5^2\right)$, true regression coefficient 0.6. We put uniform prior in $(-1, 1)$ on β. Then the true posterior

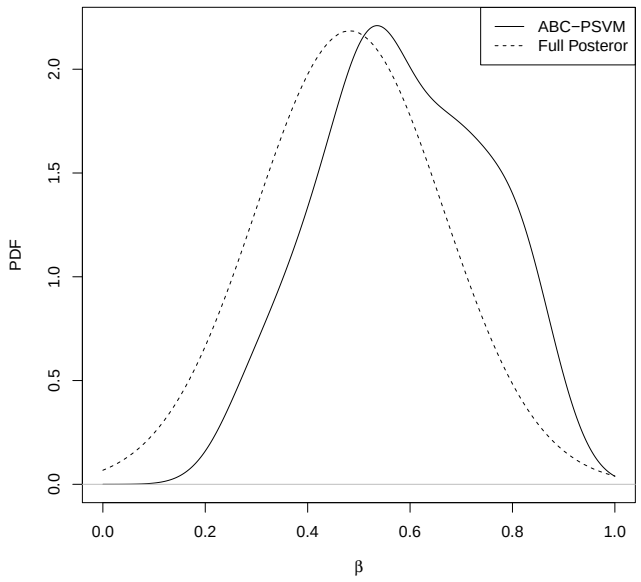

FIG 2. *ABC Density vs True Posterior Density*

distribution is $N\left(\sum_{i=1}^{99} Y_i Y_{i+1} / \left(1+\sum_{i=1}^{99} Y_i^2\right), \left(1+\sum_{i=1}^{99} Y_i^2\right)^{-1}\right)$. Now we apply our algorithm with the target dimension $d = 1$ and slicing pieces $h = 4$ with the slicing parameters y_k as quartiles. The sample size from the prior is 1000, with $k = 100/2 = 500$. Kernel κ is chosen as Gaussian kernel $\kappa(x_i, x_j) = \exp\left(-10^{-5} \times \|x_i - x_j\|^2\right)$. Then the posterior density estimated from ABC samples are plotted in Fig. 2.

The slight skewness in Fig. 2 possibly due the the small sample size of the observed data. Another interesting result of this simulation is shown in Fig. 3. There is a strong linear relationship between the estimated summary statistic and MLE

$$\hat{\beta} = \frac{\sum_{i=1}^{99} Y_i Y_{i+1}}{\sum_{i=1}^{99} Y_i^2},$$

which is one of the most efficient summary statistics based on Theorems 1 and 2. Hence, our algorithm will automatically approach the most efficient summary statistics in a nonparametric way.

4. Discussion. In this paper, we explore ABC both from theoretical and computational points of view. The theory part architects the foundation of ABC by

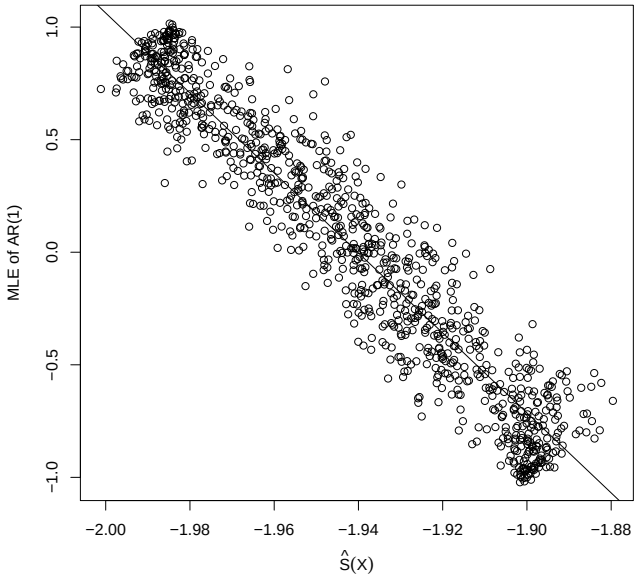

FIG 3. *Estimated Summary Statistic vs MLE*

linking asymptotic properties of statistics to that of the partial posterior. The application part innovates the algorithm by virtue of bridging selection of summary statistics and SDR. However, although the theory in Li (1992) is very powerful and may be used as a theoretical guide for our algorithm, it heavily depends on the relation (1.2) holding rigorously. We do not know whether the result from the principal support vector machine would be defunct if (1.2) is only valid in ε-sufficient way. Moreover, bringing in dimension reduction regression settings perhaps moderates the usage when there are multiple parameters of interest, and may need advance techniques such as envelope models of Su and Cook (2011; 2012).

APPENDIX A: PROOF OF THEOREM 1

PROOF.

$$\begin{aligned}
&\operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t \mid \hat{\theta}\in O\left(\theta_0,\varepsilon\right)\right\}\\
=\;&\frac{\operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t, \hat{\theta}\in O\left(\theta_0,\varepsilon\right)\right\}}{\operatorname{pr}\left\{\hat{\theta}\in O\left(\theta_0,\varepsilon\right)\right\}}\\
=\;&\frac{\int I_{\left\{\hat{\theta}\in O(\theta_0,\varepsilon)\right\}} I_{\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t\right\}} \prod_{i=1}^{n} f\left(X_i \mid \theta\right)\pi\left(\theta\right)\,\mathrm{d}X_i\,\mathrm{d}\theta}{\int I_{\left\{\hat{\theta}\in O(\theta_0,\varepsilon)\right\}} \prod_{i=1}^{n} f\left(X_i \mid \theta\right)\pi\left(\theta\right)\,\mathrm{d}X_i\,\mathrm{d}\theta}.
\end{aligned}$$

Let $\operatorname{pr}^{\infty}\left(\theta\right)$ be the probability measure on infinite independent and identically distributed sequence $X_1,\ldots,X_n,\ldots,$. Then

$$\begin{aligned}
&\operatorname{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t \mid \hat{\theta}\in O\left(\theta_0,\varepsilon\right)\right\}\\
=\;&E_{\pi(\theta)}\left(E_{\operatorname{pr}^{\infty}(\theta_0)}\left[I_{\left\{\hat{\theta}\in O(\theta_0,\varepsilon)\right\}} I_{\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t\right\}}\right.\right.\\
&\left.\left.\exp\left\{\sum_{i=1}^{n}\log\pi\left(X_i \mid \theta\right)-\log\pi\left(X_i \mid \theta_0\right)\right\}\right]\right)\\
&\Big/ E_{\pi(\theta)}\left(E_{\operatorname{pr}^{\infty}(\theta_0)}\left[I_{\left\{\hat{\theta}\in O(\theta_0,\varepsilon)\right\}}\right.\right.\\
&\left.\left.\exp\left\{\sum_{i=1}^{n}\log\pi\left(X_i \mid \theta\right)-\log\pi\left(X_i \mid \theta_0\right)\right\}\right]\right) \qquad \text{(A.1)}
\end{aligned}$$

By strong consistency of the maximum likelihood estimator, for any $\varepsilon>0$, there is an $N>0$, such that for any $n>N$, $\operatorname{pr}\left\{\hat{\theta}\in O\left(\theta_0,\varepsilon\right) \mid \theta_0\right\}=1$. Thus, we can

drop the indicator $I_{\left\{\hat{\theta}\in O(\theta_0,\varepsilon)\right\}}$ without changing the value in (A.1). We can change the order of integration. So the numerator of (A.1) is

$$\begin{aligned}
& E_{\mathrm{pr}^{\infty}(\theta_0)}\left(\left\{\prod_{i=1}^{n} f\left(X_i \mid \theta_0\right)\right\}^{-1} E_{\pi(\theta)}\left[I_{\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t\right\}}\prod_{i=1}^{n} f\left(X_i \mid \theta\right)\right]\right) \\
= \; & E_{\mathrm{pr}^{\infty}(\theta_0)}\left[\left\{\prod_{i=1}^{n} f\left(X_i \mid \theta_0\right)\right\}^{-1}\int I_{\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t\right\}}\prod_{i=1}^{n} f\left(X_i \mid \theta\right)\pi\left(\theta\right)\,\mathrm{d}\theta\right] \\
= \; & E_{\mathrm{pr}^{\infty}(\theta_0)}\left[\left\{\prod_{i=1}^{n} f\left(X_i \mid \theta_0\right)\right\}^{-1}\mathrm{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t \mid X_1,\ldots,X_n\right\}\right. \\
& \left. E_{\pi(\theta)}\left\{\prod_{i=1}^{n} f\left(X_i \mid \theta\right)\right\}\right].
\end{aligned}$$

By Bernstein–von Mises theorem,

$$\lim_{n\to\infty}\mathrm{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)\leq t \mid X_1,\ldots,X_n\right\}=\Phi\left(t\right),\ \text{a.s. pr}^{\infty}\left(\theta_0\right).$$

Hence, the result holds. □

To prove a similar result about conditioning on posterior mean, we go through similar steps as in the proof of Theorem 1. The only needed change is to prove

$$\lim_{n\to\infty}\mathrm{pr}\left[\left(n\hat{I}\right)^{1/2}\left\{\theta-E\left(\theta \mid X_1,\ldots,X_n\right)\right\}\leq t \mid X_1,\ldots,X_n\right]=\Phi\left(t\right),\ \text{a.s. pr}^{\infty}\left(\theta_0\right).$$

We know from Ghosh and Liu (2011) that with probability 1, (2.2) holds. By conditioning on $X_1,\ldots,X_n$, both posterior mean and maximum likelihood estimator are fixed numbers and

$$\lim_{n\to\infty} n^{1/2}\left\{E\left(\theta \mid X_1,\ldots,X_n\right)-\hat{\theta}\right\}=0.\ \text{a.s.}.$$

Hence, if we assume the CDF of the full posterior is continuous and asymptotically

normal, then

$$\begin{aligned}
& \left|\mathrm{pr}\left[\left(n\hat{I}\right)^{1/2}\left\{\theta - E\left(\theta \mid X_1,\ldots,X_n\right)\right\} \le t \mid X_1,\ldots,X_n\right] - \Phi\left(t\right)\right| \\
\le\ & \left|\mathrm{pr}\left[\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) + \hat{I}^{1/2}n^{1/2}\left\{E\left(\theta \mid X_1,\ldots,X_n\right) - \hat{\theta}\right\} \le t \mid X_1,\ldots,X_n\right]\right. \\
& \left. -\mathrm{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid X_1,\ldots,X_n\right\}\right| \\
& + \left|\mathrm{pr}\left\{\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) \le t \mid X_1,\ldots,X_n\right\} - \Phi\left(t\right)\right| \\
\to\ & 0, \text{ as } \left(n\to\infty\right).
\end{aligned}$$

APPENDIX B: DERIVATION OF EXAMPLES

B.1. Derivation of Example 1. First, we get the distribution of μ.

$$\mu^{r_2}\exp\left(-\mu\frac{r_2}{\hat{\mu}}\right)\left|-\frac{r_2}{\left(\hat{\mu}\right)^2}\right| \propto r_2\mu^{r_2}\exp\left(-\frac{\mu}{\hat{\mu}}r_2\right).$$

Now we sum over r_2. By definition $X_0 + r_1 - r_2 \ge 0$. Hence the distribution of $\hat{\mu}$ is proportional to

$$\sum_{r_2=0}^{X_0+r_1} r_2\left\{\mu\exp\left(-\frac{\mu}{\hat{\mu}}\right)\right\}^{r_2}.$$

Let $U = \mu\exp\left(-\mu/\hat{\mu}\right)$ and $R = X_0 + r_1 = X_0 + \hat{\lambda}T$. Let

$$\begin{aligned}
L &= \sum_{r_2=0}^{R} r_2U^{r_2} = U\frac{\mathrm{d}}{\mathrm{d}U}\sum_{r_2=0}^{R}U^{r_2} = U\frac{\mathrm{d}}{\mathrm{d}U}\left(\frac{1-U^{R+1}}{1-U}\right) \\
&= U\left(1-U\right)^{-2}\left\{1-\left(R+1\right)U^R + RU^{R+1}\right\}.
\end{aligned}$$

For fixed t, consider $\mu = \hat{\mu} + t/T^{1/2}$.

$$\begin{aligned}
\log U &= -\frac{\mu}{\hat{\mu}} + \log\mu = -1 - \frac{t}{T^{1/2}\hat{\mu}} + \log\hat{\mu} + \log\left(1+\frac{t}{T^{1/2}\hat{\mu}}\right) \\
&= \log\hat{\mu} - 1 - \frac{t^2}{2\left(\hat{\mu}\right)^2T} + o\left(T^{-1}\right).
\end{aligned}$$

Hence

$$\lim_{T\to\infty} U = \frac{\mu_0}{e}, \text{ a.s. ..}$$

Next,

$$R\log U=\left(X_0+\hat{\lambda}T\right)\left(\log\hat{\mu}-1\right)-\frac{\hat{\lambda}t^2}{2\left(\hat{\mu}\right)^2}+o\left(1\right).$$

$$L=\frac{UR}{\left(1-U\right)^2}\left[\frac{1}{R}-\frac{R+1}{R}\exp\left(R\log U\right)+U\exp\left(R\log U\right)\right]$$

Note that the density of $t=T^{1/2}\left(\mu-\hat{\mu}\right)$ is only proportional to L, hence only the terms containing t will affect the limit distribution, other terms can be omitted. Also recall that,

$$\lim_{T\to\infty}\frac{1}{R}=\lim_{T\to\infty}\frac{1}{X_0+\hat{\lambda}T}=0.$$

Thus

$$\begin{aligned}\lim_{T\to\infty}L &\propto \lim_{T\to\infty}\left\{\frac{1}{R}-\frac{R+1}{R}\exp\left(R\log U\right)+U\exp\left(R\log U\right)\right\}\\ &= \lim_{T\to\infty}\left(U-1-\frac{1}{R}\right)\exp\left\{\left(X_0+\hat{\lambda}T\right)\left(\log\hat{\mu}-1\right)-\frac{\hat{\lambda}t^2}{2\left(\hat{\mu}\right)^2}+o\left(1\right)\right\}\\ &= \lim_{T\to\infty}\left(\frac{\mu_0}{e}-1\right)\exp\left\{\left(X_0+\hat{\lambda}T\right)\left(\log\hat{\mu}-1\right)\right\}\exp\left\{-\frac{\hat{\lambda}t^2}{2\left(\hat{\mu}\right)^2}+o\left(1\right)\right\}\\ &\propto \exp\left\{-\frac{\hat{\lambda}t^2}{2\left(\hat{\mu}\right)^2}\right\}.\end{aligned}$$

B.2. Derivation of Example 2. We know $\overline{X}\sim\text{Gamma}\left(n\alpha,\beta/n\right)$, so $\tilde{\alpha}\sim\text{Gamma}\left(n\alpha,n^{-1}\right)$.

$$\begin{aligned}\pi\left(\alpha\mid\tilde{\alpha}\right) &\propto \pi\left(\tilde{\alpha}\mid\alpha\right)\pi\left(\alpha\right)=\frac{1}{\Gamma\left(n\alpha\right)n^{-n\alpha}}\left(\tilde{\alpha}\right)^{n\alpha-1}\exp\left(-n\tilde{\alpha}\right)\exp\left(-\lambda\alpha\right)\\ &\propto \frac{\left\{n\tilde{\alpha}\exp\left(-\lambda/n\right)\right\}^{n\alpha}}{\Gamma\left(n\alpha\right)}.\end{aligned}$$

Next we will show $\pi\left(n^{1/2}\left(\alpha-\tilde{\alpha}\right)/b\mid\tilde{\alpha}\right)\to N\left(0,1\right)$ a.s., for some suitable b. The PDF of t is proportional to

$$\frac{\left\{n\tilde{\alpha}\exp\left(-\lambda/n\right)\right\}^{n\left(bt/n^{1/2}+\tilde{\alpha}\right)}}{\Gamma\left\{n\left(bt/n^{1/2}+\tilde{\alpha}\right)\right\}}\frac{b}{n^{1/2}}.$$

Take logarithm, and drop all the terms not related to t, since those terms can be divided from both numerator and denominator,

$$n^{1/2}bt\log\left(n\tilde{\alpha}\right)-\lambda b\frac{t}{n^{1/2}}-\log\Gamma\left\{n\tilde{\alpha}\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)\right\}. \tag{B.1}$$

Using Stirling formula to approximate gamma function,

$$\begin{aligned}
&\log\Gamma\left\{n\tilde{\alpha}\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)\right\}\\
\approx\ &\left\{n\tilde{\alpha}\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)-\frac{1}{2}\right\}\log\left\{n\tilde{\alpha}\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)\right\}-n\tilde{\alpha}\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)+\frac{1}{2}\log 2\pi.
\end{aligned}$$

So (B.1) can be written as

$$\begin{aligned}
&n^{1/2}bt\log\left(n\tilde{\alpha}\right)-\lambda b\frac{t}{n^{1/2}}-n^{1/2}bt\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)\\
&-n^{1/2}bt\log\left(n\tilde{\alpha}\right)-\left(n\tilde{\alpha}-\frac{1}{2}\right)\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)+n^{1/2}bt\\
=\ &-n^{1/2}bt\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)-\left(n\tilde{\alpha}-\frac{1}{2}\right)\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)\\
&+n^{1/2}bt-\lambda b\frac{t}{n^{1/2}}.
\end{aligned}\tag{B.2}$$

Now we can apply Taylor expansion for term $\log\left\{1+bt/\left(n^{1/2}\tilde{\alpha}\right)\right\}$,

$$\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)=\frac{bt}{n^{1/2}\tilde{\alpha}}-\frac{b^2t^2}{2n\tilde{\alpha}^2}+\frac{b^3t^3}{3n^{3/2}\tilde{\alpha}^3}+o\left(\frac{t^3}{n^{3/2}}\right).$$

Substituting the expansion into (B.2),

$$\begin{aligned}
&-n^{1/2}bt\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)-\left(n\tilde{\alpha}-\frac{1}{2}\right)\log\left(1+\frac{bt}{n^{1/2}\tilde{\alpha}}\right)+n^{1/2}bt-\lambda b\frac{t}{n^{1/2}}\\
=\ &-n^{1/2}bt\left\{\frac{bt}{n^{1/2}\tilde{\alpha}}-\frac{b^2t^2}{2n\tilde{\alpha}^2}+\frac{b^3t^3}{3n^{3/2}\tilde{\alpha}^3}+o\left(\frac{t^3}{n^{3/2}}\right)\right\}\\
&-\left(n\tilde{\alpha}-\frac{1}{2}\right)\left\{\frac{bt}{n^{1/2}\tilde{\alpha}}-\frac{b^2t^2}{2n\tilde{\alpha}^2}+\frac{b^3t^3}{3n^{3/2}\tilde{\alpha}^3}+o\left(\frac{t^3}{n^{3/2}}\right)\right\}+n^{1/2}bt-\lambda b\frac{t}{n^{1/2}}\\
=\ &-\frac{b^2t^2}{\tilde{\alpha}}+\frac{b^3t^3}{2n^{1/2}\tilde{\alpha}^2}-o\left(\frac{t^3}{n^{1/2}}\right)-n^{1/2}bt+\frac{b^2t^2}{2\tilde{\alpha}}-\frac{b^3t^3}{2n^{1/2}\alpha^2}-o\left(\frac{t^3}{n^{1/2}}\right)\\
&+\frac{bt}{2n^{1/2}\tilde{\alpha}}-\frac{b^2t^2}{4n\tilde{\alpha}^2}+o\left(\frac{t^2}{n}\right)+n^{1/2}bt-\lambda b\frac{t}{n^{1/2}}\\
\approx\ &-\frac{b^2t^2}{2\tilde{\alpha}}.
\end{aligned}$$

If we set $b=\tilde{\alpha}^{1/2}$, then the rescaled partial posterior convergence to standard normal.

B.3. Derivation of Example 3. First we prove lemma 1.

PROOF. Assume Y has a probability density function $f(y)$. Let $y = u(x,\theta)$ be the solution of equation $x = h(y,\theta)$. Then $h(Y,\theta)$ has a probability density function

$$f(u(x,\theta))\left|\frac{\partial u(x,\theta)}{\partial x}\right|.$$

Then the posterior distribution under the uniform prior is proportional to

$$f(u(x,\theta))\left|\frac{\partial u(x,\theta)}{\partial x}\right|.$$

Now we find the probability density function of $g(Y,X)$. By assumptions, we know $y = u(x,\theta)$ is also the solution of $\theta = g(y,x)$. Hence the probability density function of $g(Y,X)$ is also

$$f(u(x,\theta))\left|\frac{\partial u(x,\theta)}{\partial x}\right|.$$

□

We know if X, Y are independent exponential random variables with mean λ, then $X - Y$ has a double exponential distribution with $\mu = 0$ and the same λ. So we know our sample Z has the same distribution as $X - Y + \mu$. So the sample mean $\overline{Z}$ has the same distribution as $\overline{X} - \overline{Y} + \mu$. It is easy to check $\overline{X}$ and $\overline{Y}$ have gamma distribution with location parameter n and scale parameter $n^{-1}\lambda$. Hence the posterior distribution of μon $\overline{Z}$ under the uniform prior has the same distribution as $\overline{Z} - \left(\overline{X} - \overline{Y}\right)$. Hence the posterior distribution $n^{1/2}\left(\mu - \overline{Z}\right)$ has the same distribution as $-n^{1/2}\left(\overline{X} - \overline{Y}\right)$ We know the characteristic function of $n^{1/2}\overline{X}$ is

$$\left\{1 - \frac{\lambda}{n} i\left(n^{1/2}t\right)\right\}^{-n},$$

So the characteristic function of $-n^{1/2}\left(\overline{X} - \overline{Y}\right)$ is

$$\left\{1 - \frac{\lambda}{n} i\left(n^{1/2}t\right)\right\}^{-n}\left\{1 - \frac{\lambda}{n} i\left(-n^{1/2}t\right)\right\}^{-n} = \left(1 + \frac{\lambda^2 t^2}{n}\right)^{-n} \to \exp\left(-\lambda^2 t^2\right).$$

Hence $n^{/1}\left(\mu - \overline{Z}\right)$ has an asymptotic normal distribution with zero mean and variance $2\lambda^2$.

APPENDIX C: PROOF OF THEOREM 2

LEMMA 2. *Under Assumptions 1, 2 and 3, for any ε, δ_1 and δ_2, there exists an N, such that for any $n \geq N$,*

$$\mathrm{pr}_{\theta_0}^{\infty}\left[\omega : \mathrm{pr}_{\omega}^{n}\left\{n^{1/2}\left|G\left(\theta,\tilde{\theta}\right) - G_1\left(\hat{\theta},\tilde{\theta}\right)\left(\theta-\tilde{\theta}\right)\right| \leq 2\varepsilon\right\} \geq 1-\delta_1\right] \geq 1-\delta_2.$$

PROOF. By Taylor expansion and Assumption 1,

$$\left|G\left(\theta,\tilde{\theta}\right) - G\left(\hat{\theta},\tilde{\theta}\right) - G_1\left(\hat{\theta},\tilde{\theta}\right)\left(\theta-\hat{\theta}\right)\right| \leq L\left(\theta-\hat{\theta}\right)^2, \tag{C.1}$$

and

$$\left|G\left(\tilde{\theta},\tilde{\theta}\right) - G\left(\hat{\theta},\tilde{\theta}\right) - G_1\left(\hat{\theta},\tilde{\theta}\right)\left(\tilde{\theta}-\hat{\theta}\right)\right| \leq L\left(\tilde{\theta}-\hat{\theta}\right)^2. \tag{C.2}$$

By posterior consistency, there exists a $\Omega_1 \subset \Omega$, $\mathrm{pr}_{\theta_0}^{\infty}(\Omega_1) = 1$, such that for any $\omega \in \Omega_1$, conditioned on $X_1(\omega), \ldots, X_n(\omega)$, θ converges in probability to θ_0. Hence conditioned on $X_1(\omega), \ldots, X_n(\omega)$, $(\theta - \theta_0) = o_{\mathrm{pr}_{\omega}^{n}}(1)$. By Bernstein–von Mises, there exists a $\Omega_2 \subset \Omega$, $\mathrm{pr}_{\theta_0}^{\infty}(\Omega_2) = 1$, such that for any $\omega \in \Omega_2$, conditioned on $X_1(\omega), \ldots, X_n(\omega)$, $\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right)$ converges in distribution to a standard normal random variable. Hence conditioned on $X_1(\omega), \ldots, X_n(\omega)$, $\left(n\hat{I}\right)^{1/2}\left(\theta-\hat{\theta}\right) = O_{\mathrm{pr}_{\omega}^{n}}(1)$. For any $\omega \in \Omega_1 \cap \Omega_2$,

$$\begin{aligned}
& \left\{n^{1/2}L\left(\theta-\hat{\theta}\right)^2 \mid X_1(\omega), \ldots, X_n(\omega)\right\} \\
= \;& L\left\{n^{1/2}\left(\theta-\hat{\theta}\right) \times \left(\theta-\hat{\theta}\right) \mid X_1(\omega), \ldots, X_n(\omega)\right\} \\
= \;& LO_{\mathrm{pr}_{\omega}^{n}}(1) \times o_{\mathrm{pr}_{\omega}^{n}}(1) = o_{\mathrm{pr}_{\omega}^{n}}(1),
\end{aligned}$$

which means for any ε and δ_1, there exists an N_1 such that

$$\mathrm{pr}_{\omega}^{n}\left[n^{1/2}L\left\{\theta-\hat{\theta}(\omega)\right\}^2 \leq \varepsilon \mid X_1(\omega), \ldots, X_n(\omega)\right] \geq 1-\delta_1.$$

By Assumptions 2, $\tilde{\theta} - \theta_0 = o_{\mathrm{pr}_{\theta_0}^{\infty}}(1)$, $n^{1/2}\left(\tilde{\theta}-\theta_0\right) = O_{\mathrm{pr}_{\theta_0}^{\infty}}(1)$, $\hat{\theta} - \theta_0 = o_{\mathrm{pr}_{\theta_0}^{\infty}}(1)$ and $n^{1/2}\left(\hat{\theta}-\theta_0\right) = O_{\mathrm{pr}_{\theta_0}^{\infty}}(1)$ .Then

$$n^{1/2}L\left(\tilde{\theta}-\hat{\theta}\right)^2 \leq 2L\left\{n^{1/2}\left(\tilde{\theta}-\theta_0\right)^2 + n^{1/2}\left(\hat{\theta}-\theta_0\right)^2\right\} = o_{\mathrm{pr}_{\theta_0}^{\infty}}(1),$$

which means for any ε and δ_2, there exists an N_2, such that for any $n \geq N_2$,

$$\mathrm{pr}_{\theta_0}^{\infty}\left[\omega : n^{1/2} L\left\{\tilde{\theta}\left(\omega\right) - \hat{\theta}\left(\omega\right)\right\}^2 \leq \varepsilon\right] \geq 1 - \delta_2.$$

Let $\Omega_\varepsilon = \left\{\omega : n^{1/2} L\left\{\tilde{\theta}\left(\omega\right) - \hat{\theta}\left(\omega\right)\right\}^2 \leq \varepsilon\right\}$. For any $\omega \in \Omega_1 \cap \Omega_2 \cap \Omega_\varepsilon$,

$$\begin{aligned}
& n^{1/2}\left|G\left(\theta, \tilde{\theta}\right) - G_1\left(\hat{\theta}, \tilde{\theta}\right)\left(\theta - \tilde{\theta}\right)\right| \\
\leq\ & n^{1/2}\left|G\left(\theta, \tilde{\theta}\right) - G\left(\hat{\theta}, \tilde{\theta}\right) - G_1\left(\hat{\theta}, \tilde{\theta}\right)\left(\theta - \hat{\theta}\right)\right| \\
& + n^{1/2}\left|G\left(\tilde{\theta}, \tilde{\theta}\right) - G\left(\hat{\theta}, \tilde{\theta}\right) - G_1\left(\hat{\theta}, \tilde{\theta}\right)\left(\tilde{\theta} - \hat{\theta}\right)\right| \\
\leq\ & n^{1/2} L\left(\theta - \hat{\theta}\right)^2 + n^{1/2} L\left(\tilde{\theta} - \hat{\theta}\right)^2 \leq 2\varepsilon,
\end{aligned}$$

with probability $1 - \delta_1$ (pr_ω^n). Also recall $\mathrm{pr}_{\theta_0}^{\infty}\left(\Omega_1 \cap \Omega_2 \cap \Omega_\varepsilon\right) \geq 1 - \delta_2$. Hence for $\omega \in \Omega_1 \cap \Omega_2 \cap \Omega_\varepsilon$ and $n \geq \max\left\{N_1, N_2\right\}$,

$$\mathrm{pr}_{\theta_0}^{\infty}\left[\omega : P_\omega^n\left\{n^{1/2}\left|G\left(\theta, \tilde{\theta}\right) - G_1\left(\hat{\theta}, \tilde{\theta}\right)\left(\theta - \tilde{\theta}\right)\right| \leq 2\varepsilon\right\} \geq 1 - \delta_1\right] \geq 1 - \delta_2.$$

□

REMARK 4. This result is weaker than the settings in posterior consistency and Bernstein–von Mises theorem. In posterior consistency, $(\theta \mid X_1, \ldots, X_n)$ converges in probability (pr_ω^n) almost surely in $\mathrm{pr}_{\theta_0}^{\infty}$. A similar comment applies to the Bernstein–von Mises theorem. However, in this lemma, the posterior distribution converges with a large probability.

REMARK 5. There is slight difference between the rescaled posterior random variable in this lemma and in

$$\begin{aligned}
(\text{C.3}) \qquad & n^{1/2}\left[\int_{\mathbb{R}} g\left(x, \tilde{\theta}\right)\pi\left(x \mid \theta\right)\mathrm{d}x - \int_{\mathbb{R}} g\left(x, \tilde{\theta}\right)\pi\left(x \mid \tilde{\theta}\right)\mathrm{d}x\right. \\
& \left. - \left\{\frac{\mathrm{d}}{\mathrm{d}\theta}\int_{\mathbb{R}} g\left(x, \tilde{\theta}\right)\pi\left(x \mid \theta\right)\mathrm{d}x\Big|_{\theta=\tilde{\theta}}\right\}\left(\theta - \tilde{\theta}\right)\right] \to 0, \text{ a.s.}.
\end{aligned}$$

The first order differential term is $G_1\left(\tilde{\theta}, \tilde{\theta}\right)$. However, since $G_1\left(\tilde{\theta}, \tilde{\theta}\right) \to G_1\left(\theta_0, \theta_0\right)$ and $G_1\left(\hat{\theta}, \tilde{\theta}\right) \to G_1\left(\theta_0, \theta_0\right)$ almost surely in $P_{\theta_0}^{\infty}$ and $n^{1/2}$ term is absorbed by $\left(\theta - \tilde{\theta}\right)$, we have proved a weak version of (C.3).

PROOF. Under the Assumptions 1, 2 and 3, we have the result from Lemma 2,

(C.4)
$$\mathrm{pr}_{\theta_0}^{\infty}\left[\omega:\mathrm{pr}_{\omega}^{n}\left\{n^{1/2}\left|G\left(\theta,\tilde{\theta}\right)-G_1\left(\hat{\theta},\tilde{\theta}\right)\left(\theta-\tilde{\theta}\right)\right|\leq 2\varepsilon\right\}\geq 1-\delta_1\right]\geq 1-\delta_2.$$

Let $\Omega_1=\left\{\omega:\mathrm{pr}_{\omega}^{n}\left\{n^{1/2}\left|G\left(\theta,\tilde{\theta}\right)-G_1\left(\hat{\theta},\tilde{\theta}\right)\left(\theta-\tilde{\theta}\right)\right|\leq 2\varepsilon\right\}\geq 1-\delta_1\right\}$, and

$$C_n=E_{\pi(\theta)}\left(E_{\mathrm{pr}^{\infty}(\theta_0)}\left[I_{\left\{\tilde{\theta}\in O(\theta_0,\varepsilon)\right\}}\exp\left\{\sum_{i=1}^{n}\log f\left(X_i\mid\theta\right)-\log f\left(X_i\mid\theta_0\right)\right\}\right]\right).$$

Now by the same technique used in the proof of Theorem 1, we have for sufficiently large N,

$$\begin{aligned}
&\left|\mathrm{pr}\left\{\frac{n^{1/2}\left(\theta-\tilde{\theta}\right)}{\tilde{V}^{1/2}}\leq t\middle|\tilde{\theta}\in O\left(\theta_0,\varepsilon\right)\right\}-\Phi\left(t\right)\right|\\
=&\left|E_{\pi(\theta)}\left(E_{\mathrm{pr}^{\infty}(\theta_0)}\left[I_{\left\{\tilde{\theta}\in O(\theta_0,\varepsilon)\right\}}I_{\left\{n^{1/2}\tilde{V}^{-1/2}\left(\theta-\tilde{\theta}\right)\leq t\right\}}\prod_{i=1}^{n}f\left(X_i\mid\theta\right)/f\left(X_i\mid\theta_0\right)\right]\right)\right.\\
&\left.\Big/E_{\pi(\theta)}\left(E_{\mathrm{pr}^{\infty}(\theta_0)}\left[I_{\left\{\tilde{\theta}\in O(\theta_0,\varepsilon)\right\}}\prod_{i=1}^{n}f\left(X_i\mid\theta\right)/f\left(X_i\mid\theta_0\right)\right]\right)-\Phi\left(t\right)\right|\\
\leq&\ C_n^{-1}E_{\pi(\theta)}\left(E_{\mathrm{pr}^{\infty}(\theta_0)}\left[I_{\Omega_1}\left|\mathrm{pr}\left\{n^{1/2}\tilde{V}^{-11/2}\left(\theta-\tilde{\theta}\right)\leq t\mid X_1,\ldots,X_n\right\}-\Phi\left(t\right)\right|\right.\right.\\
&\left.\left.\prod_{i=1}^{n}f\left(X_i\mid\theta\right)/f\left(X_i\mid\theta_0\right)\right]\right)\\
&+C_n^{-1}E_{\pi(\theta)}\left(E_{\mathrm{pr}^{\infty}(\theta_0)}\left[I_{\Omega_1^c}\left|\mathrm{pr}\left\{n^{1/2}\tilde{V}^{-11/2}\left(\theta-\tilde{\theta}\right)\leq t\mid X_1,\ldots,X_n\right\}-\Phi\left(t\right)\right|\right.\right.\\
&\left.\left.\prod_{i=1}^{n}f\left(X_i\mid\theta\right)/f\left(X_i\mid\theta_0\right)\right]\right).
\end{aligned}$$

First considering the samples within Ω_1, by (C.4), $V_0^{-1/2}\left[n^{1/2}\left\{G_1\left(\tilde{\theta},\tilde{\theta}\right)\left(\theta-\tilde{\theta}\right)-G\left(\theta,\tilde{\theta}\right)\right\}\right]$ converges in probability to 0. By Theorem 2.1 in Rivoirard et al. (2012), we have,

$$\begin{aligned}
\text{(C.5)}\qquad&\lim_{n\to\infty}\sup_{t\in\mathbb{R}}\left|\mathrm{pr}\left[n^{1/2}V_0^{-1/2}\left\{G\left(\theta,\tilde{\theta}\right)-\frac{1}{n}\sum_{i=1}^{n}g\left(X_i,\tilde{\theta}\right)\right\}\right.\right.\\
&\left.\left.\leq t\mid X_1,\ldots,X_n\right]-\Phi\left(t\right)\right|=0,\ \text{a.s.}\,\theta_0,
\end{aligned}$$

where

$$V_0 = \int_{\mathbb{R}} \left\{ g\left(x, \tilde{\theta}\right) - \int_{\mathbb{R}} g\left(y, \tilde{\theta}\right) \pi\left(y \mid \theta_0\right) \, \mathrm{d}y \right\}^2 \pi\left(x \mid \theta_0\right) \, \mathrm{d}x.$$

Hence, $n^{1/2} V_0^{-1/2} \left\{ G\left(\theta, \tilde{\theta}\right) - n^{-1} \sum_{i=1}^n g\left(X_i, \tilde{\theta}\right) \right\}$ converges in distribution to standard normal distribution. By the definition of an M-estimator, $n^{-1} \sum_{i=1}^n g\left(X_i, \tilde{\theta}\right) = 0$. Assume that for every θ, $\int_{\mathbb{R}} g\left(x, t\right) \pi\left(x \mid \theta\right) \, \mathrm{d}x = 0$ has only one solution $t = \theta$, then

$$G\left(\tilde{\theta}, \tilde{\theta}\right) = \int_{\mathbb{R}} g\left(x, \tilde{\theta}\right) \pi\left(x \mid \tilde{\theta}\right) \, \mathrm{d}x = 0.$$

Hence, by Slutsky's Theorem,

$$\begin{aligned} n^{1/2} \tilde{V}^{-1/2} \left(\theta - \tilde{\theta}\right) &= n^{1/2} V_0^{-1/2} \left\{ G\left(\theta, \tilde{\theta}\right) - \frac{1}{n} \sum_{i=1}^n g\left(X_i, \tilde{\theta}\right) \right\} \\ &\quad + V_0^{-1/2} \left[n^{1/2} \left\{ G_1\left(\tilde{\theta}, \tilde{\theta}\right) \left(\theta - \tilde{\theta}\right) - G\left(\theta, \tilde{\theta}\right) \right\} \right], \end{aligned}$$

converges in distribution to the standard normal distribution. Hence for large N,

$$\sup_{t \in \mathbb{R}} \left| \mathrm{pr}\left\{ n^{1/2} \tilde{V}^{-1/2} \left(\theta - \tilde{\theta}\right) \le t \mid X_1, \ldots, X_n \right\} - \Phi\left(t\right) \right| \le \varepsilon,$$

and

$$\begin{aligned} & C_n^{-1} E_{\pi(\theta)} \left(E_{\mathrm{pr}^\infty(\theta_0)} \left[I_{\Omega_1} \sup_{t \in \mathbb{R}} \left| \mathrm{pr}\left\{ n^{1/2} \tilde{V}^{-1/2} \left(\theta - \tilde{\theta}\right) \le t \mid X_1, \ldots, X_n \right\} - \Phi\left(t\right) \right| \right. \right. \\ & \left. \left. \prod_{i=1}^n f\left(X_i \mid \theta\right) / f\left(X_i \mid \theta_0\right) \right] \right) \\ \le \; & \varepsilon C_n^{-1} E_{\pi(\theta)} \left[E_{\mathrm{pr}^\infty(\theta_0)} \left\{ I_{\Omega_1} \prod_{i=1}^n f\left(X_i \mid \theta\right) / f\left(X_i \mid \theta_0\right) \right\} \right] = \varepsilon. \end{aligned}$$

For samples outside Ω_1. It is trivial that

$$\sup_{t \in \mathbb{R}} \left| \mathrm{pr}\left\{ n^{1/2} \tilde{V}^{-1/2} \left(\theta - \tilde{\theta}\right) \le t \mid X_1, \ldots, X_n \right\} - \Phi\left(t\right) \right| \le 2.$$

By Assumption 4 and the strong law of large numbers, and the property of the Kullback-Leibler information number

$$\begin{aligned} & \prod_{i=1}^n f\left(X_i \mid \theta\right) / f\left(X_i \mid \theta_0\right) \\ = \; & \exp\left[n \left\{ \frac{1}{n} \sum_{i=1}^n \log f\left(X_i \mid \theta\right) - \frac{1}{n} \sum_{i=1}^n \log f\left(X_i \mid \theta_0\right) \right\} \right] \le 1. \text{ a.s. } \left(P_{\theta_0}\right) \end{aligned}$$

Hence

$$C_n^{-1} E_{\pi(\theta)}\left(E_{\mathrm{pr}^\infty(\theta_0)}\left[I_{\Omega_1^c}\sup_{t\in\mathbb{R}}\left|\mathrm{pr}\left\{n^{1/2}\tilde{V}^{-1/2}\left(\theta-\tilde{\theta}\right)\le t\mid X_1,\ldots,X_n\right\}-\Phi(t)\right|\right.\right.$$
$$\left.\left.\prod_{i=1}^{n} f(X_i\mid\theta)/f(X_i\mid\theta_0)\right]\right)$$
$$\le 2C_n^{-1}E_{\pi(\theta)}\left\{E_{\mathrm{pr}^\infty(\theta_0)}\left(I_{\Omega_1^c}\right)\right\}=2\mathrm{pr}^\infty\left(\Omega_1^c\mid\theta_0\right)=2\delta_2.$$

Hence, combining (C.6) and (C.7),

$$\sup_{t\in\mathbb{R}}\left|\mathrm{pr}\left\{\frac{n^{1/2}\left(\theta-\tilde{\theta}\right)}{\tilde{V}^{1/2}}\le t\middle|\tilde{\theta}\in O(\theta_0,\varepsilon)\right\}-\Phi(t)\right|\le\varepsilon+2\delta_2,\ \text{a.s.}$$

□

APPENDIX D: PROOF OF THEOREM 3

PROOF. By Theorem 2.1 in Rivoirard et al. (2012), we have

$$\lim_{n\to\infty}\sup_{t\in\mathbb{R}}\left|\mathrm{pr}\left[n^{1/2}\mathrm{var}_{\theta_0}\left\{a^Tg(X)\right\}^{-1/2}\right.\right.$$
$$\left\{\int a^Tg(x)\,f(x\mid\theta)\,\mathrm{d}x-\frac{1}{n}\sum_{i=1}^{n}a^Tg(X_i)\right\}$$
$$\left.\left.\le t\mid X_1,\ldots,X_n\right]-\Phi(t)\right|=0,\ \text{a.s.}. \tag{D.1}$$

By Assumption 6 and Slutsky's theorem, we know $n^{-1/2}\sum_{i=1}^n g(X_i)$ has the same asymptotic distribution as S. However, by central limit theorem, $n^{-1/2}\sum_{i=1}^n g(X_i)$ has an asymptotic normal distribution with variance matrix as $\mathrm{var}_{\theta_0}\{g(X)\}$. Hence, $\mathrm{var}_{\theta_0}\{g(X)\}=\Sigma(\theta_0)=\lim_{n\to\infty}\tilde{\Sigma}$, a.s.. Hence, we can replace $\mathrm{var}_{\theta_0}\{g(X)\}$ in (D.1) by its strong consistent estimator, and get

$$\lim_{n\to\infty}\sup_{t\in\mathbb{R}}\left|\mathrm{pr}\left[n^{1/2}\left(a^T\tilde{\Sigma}a\right)^{-1/2}\left\{\int a^Tg(x)\,f(x\mid\theta)\,\mathrm{d}x-n^{-1}\sum_{i=1}^{n}a^Tg(X_i)\right\}\right.\right.$$
$$\left.\left.\le t\mid X_1,\ldots,X_n\right]-\Phi(t)\right|=0,\ \text{a.s.}.$$

By Assumption 6, we can replace $n^{1/2}\left\{n^{-1}\sum_{i=1}^{n} a^T g\left(X_i\right)\right\}$ by $n^{1/2}a^T S$, and finally obtain

$$\lim_{n\to\infty}\sup_{t\in\mathbb{R}}\left|\mathrm{pr}\left[n^{1/2}\left(a^T\tilde{\Sigma}a\right)^{-1/2}\left\{\int a^T g\left(x\right) f\left(x \mid \theta\right)\,\mathrm{d}x - a^T S\right\}\right.\right.$$
$$\left.\left.\leq t \mid X_1,\ldots,X_n\right] - \Phi\left(t\right)\right| = 0,\ \text{a.s.}.$$

The remainder of the proof uses the same argument used in the proof of Theorem 1. □

APPENDIX E: PROOF OF THEOREM 4

PROOF. First we will prove for a single slice, the support vector machine is robust by showing the normal vectors of separate hyperplanes ψ are continuous functionals of conditional probability of $\tilde{\theta}$ given $X_1,\ldots,X_n$, where $\tilde{\theta} = I_{[\theta\leq s]} - I_{[\theta>s]}$ is sliced θ. Using notations in Li, Artemiou and Li (2011), Let

$$m\left(\psi,t,X,\tilde{\theta}\right) = \psi^T\Sigma\psi + \lambda\left\{1-\tilde{\theta}\left(\psi^T X - t\right)\right\}^+,$$

be the sample version of the Lagrangian of SVM. The population normal vectors ψ are defined as the solution of the first order condition of the optimization problem in SVM,

(E.1)
$$0 = D_{(\psi,t)}E\left\{m\left(\psi,t,X,\tilde{\theta}\right)\right\} = \left(2\psi^T\Sigma,0\right)^T - \lambda E\left[\left(X,-1\right)^T\tilde{\theta}I_{\left\{1-\tilde{\theta}\left(\psi^T X-t\right)>0\right\}}\right],$$

where $D_{(\psi,t)}$ are partial derivatives over ψ and t. Let the conditional probability of $\tilde{\theta}$ given $X_1,\ldots,X_n$ be $p\left(x\right) = \mathrm{pr}\left(\tilde{\theta}=1 \mid X_1 = x_1,\ldots,X_n = x_n\right)$. We need to show that $\psi\left(p\right)$ as a functional of $p\left(x\right)$ defined by E.1 is continuous. The main theorem we rely on is Theorem 3.1.2 in Lebedev and Vorovich (2003). Next, we will check the three conditions in that theorem.

Condition (i) is trivial. For verifying Condition (ii), we first view $p\left(X\right)$ as an element from Banach space $\{p\left(x\right) : \sup_x\left|p\left(x\right)\right| < \infty\}$, with the norm $\sup_x\left|\cdot\right|$. Let

$$g\left(X,\psi,t\right) = \begin{cases} 2p\left(X\right)-1, & \psi^T X - t < 1, \\ p\left(X\right), & \psi^T X - t \geq 1, \\ p\left(X\right)-1, & \psi^T X - t \leq -1. \end{cases}$$

Then the second term in E.1, which is the only term containing $p\left(X\right)$ can be written

as

$$\begin{aligned}
& E\left[(X,-1)^T \tilde{\theta} I_{\left\{1-\tilde{\theta}(\psi^T X-t)>0\right\}}\right] \\
= \; & E_X\left((X,-1)^T E_{\tilde{\theta}|X}\left[\tilde{\theta} I_{\left\{1-(\psi X-t)\tilde{\theta}>0\right\}}\right]\right) \\
= \; & E_X\left\{(X,-1)^T g\left(X,\psi,t\right)\right\} \\
= \; & \int_{\{\psi^T X-t<1\}} (x,-1)^T \left(2p\left(x\right)-1\right) \mathrm{d}F\left(x\right) + \int_{\{\psi^T X-t\geq 1\}} (x,-1)^T p\left(x\right) \mathrm{d}F\left(x\right) \\
& + \int_{\{\psi^T X-t\leq -1\}} (x,-1)^T \left(p\left(x\right)-1\right) \mathrm{d}F\left(x\right),
\end{aligned}$$

where $F\left(x\right)$ is the CDF of marginal distribution of $X_1,\ldots,X_n$. This is a continuous linear functional map $p\left(x\right) \mapsto E\left[(X,-1)^T \tilde{\theta} I_{\left\{1-(\psi X+t)\tilde{\theta}>0\right\}}\right]$. So Condition (ii) holds. By Theorem 5 in Li, Artemiou and Li (2011), $D_{(\psi,t)} E\left\{m\left(\psi,t,X,\tilde{\theta}\right)\right\}$ can be further differentiated around the solutions. Hence, Condition (iii) holds. By Theorem 3.1.2 in Lebedev and Vorovich (2003), ψ is continuous in $p\left(x\right)$.

Let $\psi_1 = \psi\left(\mathrm{pr}\left(\tilde{\theta}=1 \mid X_1,\ldots,X_n\right)\right)$, $\psi_2 = \psi\left(\mathrm{pr}\left(\tilde{\theta}=1 \mid \Gamma_n^T\left(X_1,\ldots,X_n\right)\right)\right)$, while by Theorem 2 in Li, Artemiou and Li (2011), $\mathrm{span}\left(\psi_2\right) \subset \mathrm{span}\left(\Gamma_n\right)$. Then for any $\varepsilon>0$, there exist an N_1 and a δ such that for any $n>N_1$,

$$\begin{aligned}
& \left|\mathrm{pr}\left\{\tilde{\theta}=1 \mid \Gamma_n^T\left(X_1,\ldots,X_n\right)\right\} - \mathrm{pr}\left(\tilde{\theta}=1 \mid X_1,\ldots,X_n\right)\right| \\
= \; & \left|\mathrm{pr}\left\{\theta\leq s \mid \Gamma_n^T\left(X_1,\ldots,X_n\right)\right\} - \mathrm{pr}\left(\theta\leq s \mid X_1,\ldots,X_n\right)\right| < \delta,
\end{aligned}$$

and hence $\left|\psi_1-\psi_2\right|<\varepsilon$. By Theorem 6 in Li, Artemiou and Li (2011), the sample version of normal vectors $\hat{\psi}_1$ is weakly consistent to the population ones. Hence there exists an $N_2 \geq N_1$, such that for any $n>N_2$,

$$\begin{aligned}
& \mathrm{pr}\left(\left|\hat{\psi}_1-\psi_2\right|\geq 2\varepsilon\right) \\
\leq \; & \mathrm{pr}\left(\left\{\left|\hat{\psi}_1-\psi_1\right|\geq\varepsilon\right\}\cup\left\{\left|\psi_1-\psi_2\right|\geq\varepsilon\right\}\right) \\
= \; & \mathrm{pr}\left(\left|\hat{\psi}_1-\psi_1\right|\geq\varepsilon\right)\leq\eta.
\end{aligned}$$

Hence $\hat{\psi}_1-\psi_2 \xrightarrow{p} 0$. By Theorem 1 in Bura and Pfeiffer (2008), $\hat{\Gamma}_n$ and Γ_n being eigenvectors satisfy $\hat{\Gamma}_n-\Gamma_n \xrightarrow{p} 0$. □

ACKNOWLEDGEMENT

Ghosh's research was partially supported by an NSF Grant.

REFERENCES

BLUM, M. G., NUNES, M. A., PRANGLE, D., SISSON, S. A. et al. (2013). A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science* **28** 189–208.

BURA, E. and PFEIFFER, R. (2008). On the distribution of the left singular vectors of a random matrix and its applications. *Statistics & Probability Letters* **78** 2275–2280.

COOK, R. D. (1998). Principal Hessian directions revisited. *Journal of the American Statistical Association* **93** 84–94.

COOK, R. D. and LI, B. (2002). Dimension reduction for conditional mean in regression. *Annals of Statistics* 455–474.

COOK, R. D. and WEISBERG, S. (1991). Sliced Inverse Regression for Dimension Reduction: Comment. *Journal of the American Statistical Association* **86** pp. 328-332.

DEAN, T. A., SINGH, S. S., JASRA, A. and PETERS, G. W. (2014). Parameter estimation for hidden Markov models with intractable likelihoods. *Scandinavian Journal of Statistics* **41** 970–987.

FEARNHEAD, P. and PRANGLE, D. (2012). Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **74** 419–474.

GHOSH, M. and LIU, R. (2011). Moment matching priors. *Sankhya A* **73** 185–201.

GRELAUD, A., ROBERT, C. P., MARIN, J.-M., RODOLPHE, F., TALY, J.-F. et al. (2009). ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis* **4** 317–335.

JAAKKOLA, T. S. and JORDAN, M. I. (2000). Bayesian parameter estimation via variational methods. *Statistics and Computing* **10** 25–37.

JASRA, A., SINGH, S. S., MARTIN, J. S. and MCCOY, E. (2012). Filtering via approximate Bayesian computation. *Statistics and Computing* **22** 1223–1237.

JOHNSON, R. A. (1970). Asymptotic expansions associated with posterior distributions. *The Annals of Mathematical Statistics* **41** 851–864.

JOYCE, P. and MARJORAM, P. (2008). Approximately sufficient statistics and Bayesian computation. *Statistical applications in genetics and molecular biology* **7**.

LEBEDEV, L. and VOROVICH, I. (2003). Functional analysis in mechanics. Revised and extended translation of the Russian edition. *Springer Monographs in Mathematics. New York, NY: Springer. xi*.

LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* **86** 316–327.

LI, K.-C. (1992). On principal Hessian directions for data visualization and dimension reduction: another application of Stein's lemma. *Journal of the American Statistical Association* **87** 1025–1039.

LI, B., ARTEMIOU, A. and LI, L. (2011). Principal support vector machines for linear and nonlinear sufficient dimension reduction. *The Annals of Statistics* 3182–3210.

LI, B. and DONG, Y. (2009). Dimension reduction for nonelliptically distributed predictors. *The Annals of Statistics* 1272–1298.

LI, B., ZHA, H. and CHIAROMONTE, F. (2005). Contour regression: a general approach to dimension reduction. *Annals of statistics* 1580–1616.

MARIN, J.-M., PILLAI, N. S., ROBERT, C. P. and ROUSSEAU, J. (2014). Relevant statistics for Bayesian model choice. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76** 833–859.

MARJORAM, P., MOLITOR, J., PLAGNOL, V. and TAVARÉ, S. (2003). Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences* **100** 15324–15328.

MARTIN, R. and LIU, C. (2013). Inferential models: A framework for prior-free posterior probabilistic inference. *Journal of the American Statistical Association* **108** 301–313.

MARTIN, R. and LIU, C. (2015). Conditional inferential models: combining information for prior-free probabilistic inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **77** 195–217.

PICCHINI, U. (2014). Inference for SDE models via approximate Bayesian computation. *Journal of Computational and Graphical Statistics* **23** 1080–1100.

PRANGLE, D., FEARNHEAD, P., COX, M. P., BIGGS, P. J. and FRENCH, N. P. (2014). Semi-automatic selection of summary statistics for ABC model choice. *Statistical applications in genetics and molecular biology* **13** 67–82.

PRITCHARD, J. K., SEIELSTAD, M. T., PEREZ-LEZAUN, A. and FELDMAN, M. W. (1999). Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular biology and evolution* **16** 1791–1798.

RIVOIRARD, V., ROUSSEAU, J. et al. (2012). Bernstein–von Mises theorem for linear functionals of the density. *The Annals of Statistics* **40** 1489–1523.

ROBERT, C. and CASELLA, G. (2013). *Monte Carlo statistical methods*. Springer Science & Business Media.

RUBIN, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics* **12** 1151–1172.

SU, Z. and COOK, R. D. (2011). Partial envelopes for efficient estimation in multivariate linear regression. *Biometrika* asq063.

SU, Z. and COOK, R. D. (2012). Inner envelopes: efficient estimation in multivariate linear regression. *Biometrika* **99** 687–702.

TAVARÉ, S., BALDING, D. J., GRIFFITHS, R. C. and DONNELLY, P. (1997). Inferring coalescence times from DNA sequence data. *Genetics* **145** 505–518.

TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the american statistical association* **81** 82–86.

WILKINSON, D. J. (2011). *Stochastic modelling for systems biology*. CRC press.

WILKINSON, R. D. (2013). Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. *Statistical applications in genetics and molecular biology* **12** 129–141.

WU, H.-M. (2008). Kernel sliced inverse regression with applications to classification. *Journal of Computational and Graphical Statistics* **17**.

XIA, Y., TONG, H., LI, W. and ZHU, L.-X. (2002). An adaptive estimation of dimension reduction space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64** 363–410.

YEH, Y.-R., HUANG, S.-Y. and LEE, Y.-J. (2009). Nonlinear dimension reduction with kernel sliced inverse regression. *Knowledge and Data Engineering, IEEE Transactions on* **21** 1590–1603.

DEPARTMENT OF STATISTICS
UNIVERSITY OF FLORIDA
32611, FLORIDA, U.S.A
E-MAIL: zhongxl@stat.ufl.edu
ghoshm@stat.ufl.edu